arXiv:1405.2432v1 [stat.ML] 10 May 2014

Functional Bandits

Long Tran-Thanh* and Jia Yuan Yu†

May 13, 2014

Abstract

We introduce the functional bandit problem, where the objective is to find an arm that optimises a known functional of the unknown arm-reward distributions. These problems arise in many settings such as maximum entropy methods in natural language processing, and risk-averse decision-making, but current best-arm identification techniques fail in these domains. We propose a new approach, that combines functional estimation and arm elimination, to tackle this problem. This method achieves provably efficient performance guarantees. In addition, we illustrate this method on a number of important functionals in risk management and information theory, and refine our generic theoretical results in those cases.

Introduction

The stochastic multi-armed bandit (MAB) model consists of a slot machine with K arms (or actions), each of which delivers rewards that are independently and randomly drawn from an unknown distribution when pulled. In the optimal-arm identification problem, the aim is to find an arm with the highest expected reward value. To do so, we can pull the arms and learn (i.e., estimate) their mean rewards. That is, our goal is to distribute a finite budget of T pulls among the arms, such that at the end of the process, we can identify the optimal arm as accurately as possible. This stochastic optimisation problem models many practical applications, ranging from keyword bidding strategy optimisation in sponsored search [Amin et al., 2012], to identifying the best medicines in medical trials [Robbins, 1952], and efficient transmission channel detection in wireless communication networks [Avner, Mannor, and Shamir, 2012].

Although this MAB optimisation model is a well-studied in the online learning community, the focus is on finding the arm with the highest expected reward value [Maron and Moore, 1993, Mnih, Szepesvári, and Audibert, 2008, Audibert, Bubeck, and Munos, 2010b, Karnin, Koren, and Somekh, 2013]. However, in many applications, we are rather interested in other statistics of the

*ECS, Faculty of Physical Sciences and Engineering, University of Southampton, Southampton, United Kingdom (ltt08r@ecs.soton.ac.uk).

†IBM Research Ireland, Damastown, Dublin, Ireland (jy@osore.ca). This work is supported in part by the EU FP7 project INSIGHT under grant 318225.

arms, which can be represented as functionals of the corresponding distribution of the arm-reward values. For example, in finance and robust optimisation, notions of monetary value of risk are typically captured by risk measure functionals. In addition, telecommunication and natural language processing metrics are usually captured by information theoretic functionals (e.g., entropy and divergence). Existing optimal-arm identification techniques cannot be applied to other functionals in a straightforward way, as they exploit the fact that the expected value can be estimated in a consistent way, without any bias. This property, however, does not always hold in the nonasymptotic regime for other functionals such as entropy, divergence, or some risk measures (e.g., value-at-risk, average value-at-risk).

Against this background, we introduce a general framework, called *functional bandit optimisation*, where optimal-arm identification means finding an arm with the optimal corresponding functional value. To do so, we first propose Batch Elimination, an efficient arm elimination algorithm (i.e., optimisation method), that is suitable for identifying the best arm with a small number T of trials. The proposed algorithm is a generalised version of the Successive Elimination [Audibert, Bubeck, and Munos, 2010b] and Sequential Halving [Karnin, Koren, and Somekh, 2013] methods.

We provide generic theoretical performance guarantees for the algorithm. We refine our results in a number of scenarios with some specific and important functionals. In particular, we focus on those with applicability to risk management and information theory. Given this, for risk management, we investigate the mean-variance, value-at-risk, and average-value-at-risk functionals. Furthermore, we also study Shannon entropy functional, a widely used information theoretic metrics.

The following are our main contributions. We start with introducing the problem of optimal-arm identification with functional bandits, a generalised framework for the existing best arm identification model. We then propose Batch Elimination, which can be regarded as a generalised version of many existing optimisation methods. We also provide rigorous theoretical performance analysis for the algorithm. In the following sections, we refine our results to a number of practical risk management and information theoretic functionals. The last section concludes with questions and discussions.

Related work

Multiarmed bandit problems have been studied in a variety of settings, including the Markovian (rested and restless), stochastic and adversarial settings. For surveys on bandit problems, we refer the reader to [Gittins, Glazebrook, and Weber, 2011, Cesa-Bianchi and Lugosi, 2006]. Two types of results are found in the literature: results on the average regret (i.e., in a regret-minimization setting [Lai and Robbins, 1985]) and results on the sample complexity (i.e., in a pure-exploration setting [Even-Dar, Mannor, and Mansour, 2002, Domingo, Gavaldá, and Watanabe, 2002, Bubeck, Munos, and Stoltz, 2011]). Our work is of the second type. It is re-

lated to work on sample complexity of bandit arm-selection [Even-Dar, Mannor, and Mansour, 2002, Kalyana et al., 2012], which is also known as pure exploration or best-arm identification [Audibert, Bubeck, and Munos, 2010a, Gabillon et al., 2012].

In the Markovian setting, [Denardo, Park, and Rothblum, 2007, Chancelier, Lara, and de Palma, 2009] consider a one-armed bandit problem in the setting of Gittins indices and model risk with concave utility functions. In the stochastic setting, the notion of risk has been limited to empirical variance [Audibert, Munos, and Szepesvári, 2009, Sani, Lazaric, and Munos, 2012] and risk measures [Yu and Nikolova, 2013]. In [Audibert, Munos, and Szepesvári, 2009, Sani, Lazaric, and Munos, 2012], the functionals assign real values to the decision-maker's *policies* (*i.e.*, confidence-bound algorithms) and guarantees are given for the *regret* in retrospect. As in [Yu and Nikolova, 2013], our functionals assign a real value to random variables, *i.e.*, *rewards* of individual arms. This is more in line with the risk notions of the finance and optimization literature.

The Functional Bandit Model

Our bandit model consists of K arms. By pulling a particular arm $i \in [1, \ldots, K]$, we receive a reward X_i drawn from an unknown stationary distribution F_i (i.e., repeatedly pulling the same arm results in generating a sequence of i.i.d. random variables). Suppose $G(.)$ is a functional of F_i, and we denote the its value for arm i as $G_i = G(F_i)$. Our goal is to identify the arm with the best (e.g., highest, or lowest) functional value. For the sake of simplicity, we assume that the highest functional is the best. That is, we aim to find $i^* = \arg\max_i G_i$. However, as F_i are initially unknown, we aim to achieve this goal by using an arm pulling policy which works as follows. For $t = 1, \ldots, T$ finite number of time steps, at each t, the policy chooses an arm $i(t)$ to pull, and observes the received reward. At the end of T, the policy chooses an arm, denoted as the random variable $i^+(T)$, which it believes has the best functional value. The regret of an arm-i is defined as $\gamma_i \triangleq G_{i^*} - G_i$. The expected regret of the policy is:

$$r(T) = \mathbb{E}[G_{i^*} - G_{i^+(T)}] = \mathbb{E}[\gamma_{i^+(T)}].$$

In addition, let

$$e_r(T) = \mathbb{P}(i^+(T) \neq i^*)$$

denote the probability that we recommend a suboptimal arm after T samples. Our goal is then to find a policy that achieves minimal regret and the recommendation error. In what follows, we describe a generic algorithm that is designed to efficiently identify the best arm.

The Batch Elimination Algorithm

We now turn to the description of Batch Elimination, our arm elimination algorithm. Its pseudo code is depicted in Algorithm 1. For a given integer $L > 0$, let $\{x_1, \ldots, x_L\}$ be a sequence of L non-negative integers such that $\sum_{l=1}^{L} x_l = K-1$.

Algorithm 1 The Batch Elimination Algorithm

Inputs: T, $\{x_1, \ldots, x_L\}$.
Initialize: $S_1 = \{1, \ldots, K\}$.
for $l = 1$ **to** L **do**
  pull each remaining arm $i \in S_m$ for $\lfloor T/H \rfloor$ times,
  use estimator $\widehat{G}_i$ to estimate G_i,
  eliminate the weakest x_l arms from S_l to obtain S_{l+1}.
end for
Output: $i^+(T)$, which is the sole arm in S_L.

The Batch Elimination algorithm runs over L rounds. Within each round $l = 1, \ldots, L$, it maintains a set S_l of remaining arms, and it pulls each of the arms within this set $\lfloor T/H \rfloor$ times, where the value of H is defined later. It then uses the corresponding reward-samples to update the functional estimate $\widehat{G}_i$ of G_i of each remaining arm. Here, we assume that we have access to an estimator for each functional G_i, which can be calculated from the samples drawn from F_i. This estimator must have a property that will be defined later. Finally, we eliminate the weakest x_m arms (i.e., those with the lowest estimates for their functionals) and proceeds to the next round.

Observe that since $\sum_{l=1}^{L} x_l = K - 1$, the above algorithm repeats until one arm remains. Our algorithm can be regarded as a generalised version of the Successive Rejects [Audibert, Bubeck, and Munos, 2010b] and Sequential Halving [Karnin, Koren, and Somekh, 2013] algorithms, which are designed for identifying the arm with highest *mean value*. In fact, by setting $L = K - 1$ and $x_1 = x_2 = \cdots = x_L = 1$, we get the Successive Rejects method, where we only eliminate the weakest remaining arm at each round. On the other hand, by setting $L = \lceil \log_2 K \rceil$, $x_1 = \lfloor K/2 \rfloor$ and $x_l = \lfloor (K - \sum_{j=1}^{l-1} x_j)/2 \rfloor$ for $1 < l \leqslant L$, we get the Sequential Halving algorithm, where we eliminate the weaker half of the remaining arms at each round.

It is remained to set the value of H. Recall that at each round, we pull each arm $\lfloor T/H \rfloor$ times. Given this, we have to choose the value of H such that the total number of pulls within the algorithm does not exceed T. Let

$$H = LK - \sum_{l=1}^{L} x_l (L - l). \tag{1}$$

We show that by doing so, we can guarantee that our algorithm does not pull more than T arms. Indeed, observe that at each round l, the total number of pulls is $|S_l| \lfloor T/H \rfloor$. Hence, the total number of pulls is

$$\sum_{l=1}^{L} |S_l| \Big\lfloor \frac{T}{H} \Big\rfloor \leqslant \sum_{l=1}^{L} \Big(K - \sum_{j=1}^{l-1} x_j \Big) \frac{T}{H} = H \frac{T}{H} = T.$$

Performance Analysis

Given the description of our arm elimination algorithm, we now turn to the investigate its performance. To do so, we first define when an estimator is considered to be sufficiently good. Let $\widehat{G}(n)$ denote the estimate value of $G(F)$ by using estimator $\widehat{G}$ on n i.i.d. samples from F. We say that:

Definition 1 (Q-efficient estimator)**.** An estimator $\widehat{G}$ of functional G is *Q-efficient* if there exists a function $Q(n, x) : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, monotone decreasing in both n and x, such that for every number of sample $n > 0$ and real number $x > 0$, we have

$$\mathbb{P}(\widehat{G}(n) - G \geqslant x) \leqslant Q(n, x)$$
$$\text{and } \mathbb{P}(G - \widehat{G}(n) \geqslant x) \leqslant Q(n, x)$$

A Q-efficient estimator has the property that the probability that the one-sided estimation error exceeds x is bounded by $Q(n, x)$. Intuitively, this property guarantees that the estimate $\widehat{G}$ is not too far away from the true value. We demonstrate in the subsequent sections that many practical functionals have Q-efficient estimators.

Next, we turn to the analysis of the Batch Elimination algorithm. First, we bound the probability of obtaining a wrong ordering of two arms from Q-efficient estimators.

Lemma 1 (Wrong order)**.** *Consider a fixed round l. Suppose that $\widehat{G}_i$ is a Q-efficient estimator for each $i \in S_l$ and that we already have n samples from each remaining arm. In addition, suppose that i^* has not been eliminated before this round. Then, for all $i \in S_l$ and $i \neq i^*$, we have:*

$$\mathbb{P}(\widehat{G}_i(n) \geqslant \widehat{G}_{i^*}(n)) \leqslant 2Q\Big(n, \frac{\gamma_i}{2}\Big)$$

Next, we bound the probability of eliminating the optimal arm i^*.

Lemma 2 (Eliminating optimal arm)**.** *Suppose that the assumptions of Lemma 1 hold. Let $d \triangleq \min_{i \neq i^*} \gamma_i$. Suppose that the best arm i^* is not eliminated until round l. The probability that the optimal arm is eliminated at the end of round l is bounded as follows:*

$$\mathbb{P}(i^* \notin S_{l+1} \mid i^* \in S_l) \leqslant 2(|S_l| - x_l)Q\Big(n, \frac{d}{2}\Big).$$

Our main result can be stated as follows.

Theorem 3 (Main result)**.** *Suppose that the assumptions of Lemma 2 hold. We have the following upper bound for the recommendation error:*

$$e_r(T) \leqslant 2\Big(H - K + 1\Big)Q\Big(\frac{T - H}{H}, \frac{d}{2}\Big).$$

Alternatively, the following corollary provides a minimal number of pulls that can guarantee a high probabilistic success.

Corollary 4 (Sample complexity)**.** *Let $0 < \delta < 1$ and let $Q^{-1}_{d/2}(.)$ denote the inverse function of $Q_{d/2}(n) = Q\Big(n, \frac{d}{2}\Big)$. If Q is strictly monotone in n, then if*

$$T \geqslant HQ^{-1}_{d/2}\Big(\frac{\delta}{2H - 2K + 2}\Big) + H,$$

our success probability (i.e., correctly recommending i^*) is at least $(1 - \delta)$.

Proof sketch of Theorem 3. Let $e_p(l)$ denote the probability that the best arm is not eliminated until round l but will be dropped out at this round. Given this, we have:

$$\begin{aligned} e_r(T) &\leqslant \sum_{j=1}^{L} e_p(j) \leqslant \sum_{j=1}^{L} 2(|S_j| - x_j)Q\Big(l\frac{T-H}{H}, \frac{d}{2}\Big) \\ &\leqslant 2Q\Big(\frac{T-H}{H}, \frac{d}{2}\Big)\sum_{j=1}^{L}\Big(K - \sum_{r=1}^{l-1} x_r - x_j\Big) \\ &\leqslant 2Q\Big(\frac{T-H}{H}, \frac{d}{2}\Big)\Big(H - K + 1\Big) \end{aligned}$$

which concludes proof. □

Proof of Corollary 4. Note that since $Q_{d/2}$ is strictly monotone decreasing in n, $Q^{-1}_{d/2}$ always exists. In addition, we have:

$$\begin{aligned} e_r(T) &\leqslant 2\Big(H - K + 1\Big)Q\Big(\frac{T-H}{H}, \frac{d}{2}\Big) \\ &\leqslant 2\Big(H - K + 1\Big)Q\Big(\frac{T_0-H}{H}, \frac{d}{2}\Big) \end{aligned}$$

where $T_0 = HQ^{-1}_{d/2}\Big(\frac{\delta}{2H-2K+2}\Big) + H$. Simplifying the last term we obtain $e_r(T) \leqslant \delta$. □

Ou main result, Theorem 3, also implies the following corollaries.

Corollary 5 (Regret bound)**.** *We have the following upper bound for the expected regret*

$$r(T) \leqslant 2\gamma_{\max}\Big(H - K + 1\Big)Q\Big(\frac{T-H}{H}, \frac{d}{2}\Big)$$

where $\gamma_{\max} = \max_{j \neq i^*} \gamma_j$.

Corollary 6 (PAC regret bound)**.** *Let the algorithm output after T time steps be $i^+(T)$, then for any $0 < \delta < 1$, we have*

$$G_{i^*} - G_{i^+(T)} \leqslant \frac{2\gamma_{\max}}{\delta}\Big(H - K + 1\Big)Q\Big(\frac{T-H}{H}, \frac{d}{2}\Big)$$

with at least $(1 - \delta)$ probability.

These results provide the upper bound for the algorithm's regret and probably approximately correct (PAC) regret, respectively. The latter is a regret that holds with high probability, in contrast to the former, which always holds.

Proof sketch of Corollary 5. By definition, we have

$$\begin{aligned} r(T) &= \sum_{j \neq i^*} \gamma_j \mathbb{P}\Big(I^+ = j\Big) \leqslant \gamma_{\max} \sum_{j \neq i^*} \mathbb{P}\Big(I^+ = j\Big) \\ &\leqslant \gamma_{\max} e_r(T) \leqslant 2\gamma_{\max}\Big(H - K + 1\Big) Q\Big(\frac{T-H}{H}, \frac{d}{2}\Big) \end{aligned}$$

□

Proof sketch of Corollary 6. Let $\varepsilon > 0$ be an arbitrary number. From Markov's inequality, we have

$$\mathbb{P}\Big(G_{i^*} - G_{i^+(T)} \leqslant \varepsilon\Big) \geqslant 1 - \frac{\mathbb{E}[G_{i^*} - G_{i^+(T)}]}{\varepsilon} \geqslant 1 - \frac{r(T)}{\varepsilon}$$

This also holds for $\varepsilon = r(T)/\delta$ where $0 < \delta < 1$. An application of Corollary 5 concludes the proof. □

As our results hold for a large class of functionals, they are not comparable with the existing optimal arm identification results in general, as the latter are only designed for the expected value functional. However, by applying our results to the case of expected values, we can still compare them against the state of the art. In what follows, we will make a comparison between our performance guarantees and the existing bounds of Audibert, Bubeck, and Munos [2010b], and Karnin, Koren, and Somekh [2013], respectively.

In particular, we can show that the expected value functional is Q-efficient with $Q(n, x) = \exp\Big\{ - nx^2 \Big\}$. Indeed, this can be proved by using Hoeffding's inequality. Now, by setting the values of L and x_l to be $L = \lceil \log_2 K \rceil$, $x_1 = \Big\lfloor K/2 \Big\rfloor$ and $x_l = \Big\lfloor (K - \sum_{j=1}^{l-1} x_j)/2 \Big\rfloor$ for $1 < l \leqslant L$, we can prove that our algorithm has

$$(K + \log_2 K) \exp\Big\{ - \frac{(T - 2K)d^2}{8K} \Big\}$$

upper bound for the recommendation error probability $e_r(T)$ when applying our results to the expected value case in a straightforward way (i.e., without using further technical refinements). Note that the upper bound of the same error probability is

$$\frac{K(K-1)}{2} \exp\Big\{ - O\Big(\frac{(T-K)d^2}{K \log_2 K}\Big) \Big\}$$

in [Audibert, Bubeck, and Munos, 2010b], and

$$3 \log_2 K \exp\Big\{ - O\Big(\frac{Td^2}{8K \log_2 K}\Big) \Big\}$$

in [Karnin, Koren, and Somekh, 2013], respectively[1]. Roughly speaking, our error bound has better constant coefficient (i.e., $K + \log_2 K$), compared to that of Audibert, Bubeck, and Munos [2010b], but weaker than that of in [Karnin, Koren, and Somekh, 2013]. On the other hand, it outperforms both existing bounds from the aspect of the exponential term. In fact, the exponential term of our bound is $\exp\{-O(\frac{T}{K})\}$, while the others have $\exp\{-O(\frac{T}{K \log_2 K})\}$. This implies that for sufficiently large values of K, our regret bound is more efficient than the existing ones. Similar results can be obtained for the comparison of regret and PAC regret bounds. It is worth to mention that by using some elementary algebra, we can improve the error bound of Batch Elimination for expected values to

$$C \log_2 K \exp\Big\{ - O\Big(\frac{T d^2}{K}\Big)\Big\}$$

for some constant $C > 0$. However, due to space limitations, we omit the details.

Recall that the abovementioned analyses rely on the assumption that the investigated functional has a Q-efficient estimator, which might not always hold. However, in the next sections, we will demonstrate that many practical functionals have this property (i.e., there is a Q-efficient estimator for such functionals). In particular, we derive specific Q functions for each of these functionals, and we refine the regret bound results, described in Theorem 3 and Corollaries 5 and 6, tailoring to each corresponding functional.

Risk Functionals

In this section, we consider three specific instances of functionals widely used as risk measures in decision-making. They are the mean-variance risk, the value-at-risk, and the average value-at-risk.

Mean-Variance

In this section, we consider the mean-variance objective functional

$$G_i^{M,\lambda} = -\mu(i) + \lambda \sigma^2(i),$$

where $\mu(i)$ and $\sigma^2(i)$ denote the mean and variance of arm i. The mean-variance risk measure has been used in risk-averse problem formulations in a variety of applications in finance and reinforcement learning Markowitz [1952], Mannor and Tsitsiklis [2011], Sani, Lazaric, and Munos [2012].

Let λ be given and fixed. We assume that we are given N samples $X_1^i, \ldots, X_N^i$ for every arm i. We employ the following unbiased estimate

$$\widehat{G}_i^{M,\lambda} = -\widehat{\mu}(i) + \lambda \widehat{\sigma}^2(i),$$

where $\widehat{\mu}(i)$ and $\widehat{\sigma}^2(i)$ denote the sample-mean and the unbiased variance estimator $\widehat{\sigma}^2(i) = \frac{1}{N-1} \sum_{k=1}^{N} (X_k^i - \widehat{\mu}(i))^2$.

[1] In these bounds, the expression in the exponential term is in fact $O(-T/(H_2 \log_2 K))$ with $H_2 = \max_{i \neq i^*} i/\gamma_i^2$. However, we can roughly estimate $\frac{1}{H_2}$ with $O(d^2/K)$.

Theorem 7 (MV PAC bound)**.** *Suppose that there exist A, B such that $\mathbb{P}(X_t^i \in [A,B]) = 1$ for all i. Let i^* denote the best arm with respect to the functional $G^{M,\lambda}$. Consider a fixed round l and N samples for each remaining arm. Then, for all $i = 1, 2 \ldots$, such that $i \neq i^*$ and is still not eliminated at round l, we have the following bound:*

$$\begin{aligned}\mathbb{P}(\widehat{G}_i^{M,\lambda}(N) \geqslant \widehat{G}_{i^*}^{M,\lambda}(N)) &\leqslant 2\exp\left(-\frac{N\gamma_i^2}{8(B-A)^2}\right) \\ &+ 2\exp\left(-\frac{N(\frac{N-1}{N}\gamma_i/\lambda)^2}{8(B-A)^4}\right).\end{aligned}$$

A straightforward application of this theorem to Lemma 2 and Theorem 3 implies the following:

Corollary 8. *For the mean-variance functional case, our algorithm has the following upper bound for the recommendation error $e_r(T)$:*

$$\begin{aligned}e_r(T) &\leqslant 2\Big(H-K+1\Big)\exp\left(-\frac{(T-H)\gamma_i^2}{8H(B-A)^2}\right) \\ &+ 2\Big(H-K+1\Big)\exp\left(-\frac{(T-2H)^2(\gamma_i/\lambda)^2}{8(T-H)(B-A)^4}\right)\end{aligned}$$

We can also derive regret and PAC regret bounds for this case, similarly to Corollaries 5 and 6. However, due to the space limitations, we leave this to the reader.

Value-at-Risk

Let λ be given and fixed. In this section, we consider the value-at-risk, for every arm i:

$$G_i^{V,\lambda} = \mathrm{V@R}_\lambda(X_1^i) = -q_i(\lambda),$$

where q_i is the right-continuous quantile function[2] of X_1^i.

Suppose that up to time T, each arm is sampled N times. Let $X_1^i, \ldots, X_N^i$ denote the sequence of rewards generated by arm i. Let $X_{(1)}^i \leqslant \ldots \leqslant X_{(N)}^i$ denote a reordering of the random variables $\{X_1^i, \ldots, X_N^i\}$, where $X_{(k)}^i$ is the k-th order statistic of the sequence $\{X_1^i, \ldots, X_N^i\}$. We consider the following V@R estimators for all i:

$$\widehat{G}_i^{V,\lambda}(N) \triangleq -X_{(\lceil \lambda N \rceil)}^i,$$

where $X_{(\lceil \lambda N \rceil)}^i$ is a λ-quantile estimator[3].

[2] Formally, $q_i(\lambda) = \inf\{x \in \mathbb{R} : F_i(x) > \lambda\}$, where F_i is the distribution function of X_1^i.

[3] With slight modifications, we can derive similar results with other quantile estimators, such as the Wilks estimator.

Assumption 1 (Differentiable reward density)**.** For each arm i, the reward probability density functions d_i are continuously differentiable.

The following theorem from the theory of order statistics establishes the convergence of the quantile estimator.

Theorem 9. *[David and Nagaraja, 2003] Suppose that Assumption 1 holds. Let d_i denote the probability density function of arm i's rewards, and d_i' denote the derivative of d_i. There exist constants $C_1, C_2 \geqslant 0$ and scalars V_N^i and W_N^i*

$$\begin{aligned} \left|V_N^i\right| &\leqslant \left|\frac{\lambda(1-\lambda)d_i'(q_i(\lambda))}{2(N+2)d_i^3(q_i(\lambda))}\right| + C_1/N^2, \\ W_N^i &\leqslant \frac{\lambda(1-\lambda)}{(N+2)d_i^2(q_i(\lambda))} + C_2/N^2 \end{aligned}$$

such that

$$\begin{aligned} \mathbb{E}X^i_{(\lceil \lambda N \rceil)} &= q_i(\lambda) + V_N^i, \\ \mathbb{V}\, X^i_{(\lceil \lambda N \rceil)} &= \mathbb{E}(X^i_{(\lceil \lambda N \rceil)} - \mathbb{E}X^i_{(\lceil \lambda N \rceil)})^2 = W_N^i. \end{aligned}$$

The following theorem uses a result from order statistics to derive a PAC sample complexity bound on our estimator.

Theorem 10 (V@R estimation error)**.** *Suppose that Assumption 1 holds. Suppose that the number of samples of each arm is N. Then, for every arm i, we have:*

$$\mathbb{P}\left(\left|\widehat{G}_i^{V,\lambda}(N) - G_i^{V,\lambda}\right| > \varepsilon\right) \leqslant \frac{W_N^i}{(\varepsilon - \left|V_N^i\right|)^2}.$$

This result implies the following statement:

Lemma 11. *Suppose i^* has not been eliminated until round l. Let N denote the total number of pulls per remaining arm. Then, for all $i \in S_l$ and $i \neq i^*$, we have:*

$$\mathbb{P}(\widehat{G}_i^{V,\lambda}(N) \geqslant \widehat{G}_{i^*}^{V,\lambda}(N)) \leqslant 2\frac{W_N^i}{(\frac{d}{2} - \left|V_N^i\right|)^2}$$

By applying this to Lemma 2 and Theorem 3, we can refine the upper bound for the recommendation error probability $e_r(T)$ of our algorithm as follows:

Corollary 12. *The recommendation error for the value-at-risk functional case has the following upper bound:*

$$e_r(T) \leqslant 4\Big(H - K + 1\Big)\frac{W}{(\frac{d}{2} - |V|)^2}$$

where $W = \max_{i \neq i^*} W^i_{\lfloor \frac{T}{H} \rfloor}$ and $V = \max_{i \neq i^*} V^i_{\lfloor \frac{T}{H} \rfloor}$.

We can also derive regret and PAC regret bounds for this case. However, due to the space limitations, we omit these steps.

Average Value-at-Risk

Modern approaches to risk measures advocate the use of convex risk measures, which capture the fact that diversification helps reduce risk. In this section, we consider only one instance of convex risk measures: the average value-at-risk. Nonetheless, it can be shown that an important subset of convex risk measures (*i.e.*, those continuous from above, law invariant, and coherent) can be expressed as an integral of the AV@R (cf. Schied [2006]). Guarantees can be obtained for those risk measures by using the approach of this section.

The AV@R has the following two equivalent definitions—first, as an integral of V@R:

$$G_i^{A,\lambda} = \mathrm{AV@R}_\lambda(X_1^i) = \frac{1}{\lambda}\int_0^\lambda \mathrm{V@R}_\phi(X_1^i)\mathrm{d}\phi,$$

and second, as a maximum over a set of distributions: $\rho_\lambda^A(X) = \max_{Q\in\mathcal{Q}_\lambda(\mathbb{P})} -\mathbb{E}_Q X$, where $\mathcal{Q}_\lambda(\mathbb{P})$ is the set of probability measures $\{Q : \frac{\mathrm{d}Q}{\mathrm{d}\mathbb{P}} \leqslant 1/\lambda\}$. Depending on the choice of definition, we can estimate the AV@R either via quantile estimation or density estimation. In this section, we adopt the first definition and introduce the following AV@R estimator of Yu and Nikolova [2013]

$$\widehat{G}_i^{A,\lambda}(N) \triangleq -\frac{1}{\lambda}\left(\sum_{j=0}^{\lfloor\lambda N\rfloor-1}\frac{1}{N}X_{(j+1)}^i + \left(\lambda - \frac{\lfloor\lambda N\rfloor}{N}\right)X_{(\lceil\lambda N\rceil)}^i\right)$$

which is a Riemann sum of V@R estimators. Observe that it is computationally more efficient than that of Brown [2007].

Theorem 13 (AV@R sample complexity)**.** *Suppose that the assumptions of Theorem 10 hold. Suppose that the rewards are bounded such that* $\left|X_t^i\right| \leqslant M$ *almost surely, for every arm* i *and time* t*. In addition, we assume that there exist* D *and* D' *such that* $d_i(z) \leqslant D$ *and* $d_i'(z) \leqslant D'$ *for all* $z \in R$ *and all* i*, and that*

$$N \geqslant \max\left\{\frac{32\lambda'(N)M^2}{\varepsilon^2\lambda^2}\log(2/\delta), \frac{(1/6)D'/D^3 + 2C_1\lambda'}{\varepsilon\lambda}, 2\right\},$$

where $\lambda'(N)$ *denotes the smallest real number greater than* λ *such that* $N\lambda'(N)$ *is an integer. Then, we have, for every arm* i*,*

$$\mathbb{P}\left(\left|G_i^{A,\lambda} - \widehat{G}_i^{A,\lambda}(N)\right| > \varepsilon\right) \leqslant \delta.$$

This implies the following statement.

Corollary 14. *Using the notation from Theorem 13, suppose that*

$$N \geqslant \max\left\{\frac{(1/6)D'/D^3 + 2C_1\lambda'(N)}{\varepsilon\lambda}, 2\right\}.$$

Given this, we have

$$\mathbb{P}\left(\left|G_i^{A,\lambda} - \widehat{G}_i^{A,\lambda}(N)\right| > \varepsilon\right) \leqslant 2\exp\left\{-\frac{N\varepsilon^2\lambda^2}{32\lambda'(N)M^2}\right\}.$$

From this corollary, we can also derive the following.

Corollary 15. *Using the notation from Theorem 13, suppose that*

$$\Big\lfloor \frac{T}{H} \Big\rfloor \geqslant \max\left\{ \frac{(1/6)D'/D^3 + 2C_1\lambda'(N)}{\varepsilon\lambda}, 2 \right\}.$$

Given this, the recommendation error for the average value-at-risk functional case has the following upper bound:

$$e_r(T) \leqslant 4\Big(H - K + 1\Big) \exp\Big\{ - \frac{(T-H)\varepsilon^2\lambda^2}{32H\lambda'(\lfloor \frac{T}{H} \rfloor)M^2} \Big\}.$$

Information Theoretic Functionals

Optimal-arm identification with information theoretic functionals are quite common in natural language processing applications. For example, it is desirable to find an arm that maximizes the entropy in [Berger, Pietra, and Pietra, 1996]. Given this, this section investigates functional bandit optimisation with entropy and its other counterparts. To do so, we first describe the required conditions that we focus on within this section. We then provide refined error probability bounds.

Assumption 2 (Discrete random variables)**.** Suppose that the rewards $X_1^i, \ldots, X_N^i$ are i.i.d. , take values in a countable set V, and with distribution F_i.

Let G_i^H denote the entropy of arm i:

$$G_i^H = H(F_i) = \sum_{v \in V} F_i(v) \log_2 F_i(v).$$

Let $F_i(N)$ denote the empirical frequency of realisations in the samples $X_1^i, \ldots, X_N^i$. A number of consistent estimators exist for the entropy: e.g., plug-in, matching-length [Antos and Kontoyiannis, 2001], and nearest-neighbour estimators. We consider the k-nearest neighbour entropy estimator of [Sricharan, Raich, and Hero, 2011].

Theorem 16. *[Sricharan, Raich, and Hero, 2011] Suppose that $X_1^i, X_2^i, \ldots$ are d-dimensional i.i.d. random variables and that the reward distribution of arm j admits a density function with bounded support. Let c_1, c_2, c_4, c_5 denote constants that depend on F_i and G_i^H only. The k-nearest neighbor entropy estimator with parameter M satisfies:*

$$\begin{aligned} V_N^i &= \mathbb{E}\left| \widehat{G}_i^H(N) - G_i^H \right| \\ &= c_1 \left(\frac{k}{M} \right)^{1/d} + c_2/k + o(1/k + (k/M)^{1/d}), \\ W_N^i &= \mathbb{E}[\widehat{G}_i^H(N) - \mathbb{E}\widehat{G}_i^H(N)]^2 \\ &= c_4/N + c_5/M + o(1/M + 1/N). \end{aligned}$$

Theorem 17 (Entropy estimation error)**.** *Suppose that the assumptions of Theorem 16 hold. We have*

$$\mathbb{P}(\left|\widehat{G}_i^H(N) - G_i^H\right| > \varepsilon) \leqslant \frac{W_N^i}{(\varepsilon - \left|V_N^i\right|)^2}.$$

The proof is similar to the proof of Theorem 10, and thus, is omitted. In addition, we have the following result, similar to the case of value-at-risk.

Corollary 18. *Suppose that the assumptions of Theorem 16 hold. Given this, the recommendation error of our algorithm for the entropy case has the following upper bound:*

$$e_r(T) \leqslant 4\Big(H - K + 1\Big)\frac{W}{(\frac{d}{2} - |V|)^2}$$

where $W = \max_{i \neq i^*} W^i_{\lfloor \frac{T}{H} \rfloor}$ *and* $V = \max_{i \neq i^*} V^i_{\lfloor \frac{T}{H} \rfloor}$.

Remark 1 (Other information theoretic functionals). A similar result holds for Rényi entropy—a generalization of the notion of Shannon entropy, cf. [Sricharan, Raich, and Hero, 2011] and [Pál, Poczós, and Szepesvári, 2010], divergence, and mutual information.

Conclusions

In this paper we introduced the functional bandit optimisation problem, where the goal is to find an arm with the optimal value of a predefined functional of the arm-reward distributions. To tackle this problem, we proposed Batch Elimination, an algorithm that combines efficient functional estimation with arm elimination. In particular, assuming that there exists a Q-efficient estimator of the functional, we run a number of arm pulling rounds, and eliminate a certain number of weakest remaining arms. The algorithm stops when there is only one arm left. We analysed the performance of the algorithm by providing theoretical guarantees on its recommendation error, regret, and PAC regret, respectively. We also refined our results in a number of cases where we use risk management and information theoretic functionals.

The most trivial way to extend our results to other functionals is via "plug-in" functional estimators. In addition, it is also feasible to modify our algorithm to output a ranking of arms. Note that when there are multiple arms that are approximately equally good, we may want to extend our analysis to give probabilities of returning the j-th best arm.

Recall that our algorithm requires the knowledge of the total number of pulls T in advance. However, this is not always the case in many applications. Given this, it is a desirable goal to extend our algorithm to run in an online fashion, without requiring the time horizon T as an input.

References

Amin, K.; Kearns, M.; Key, P.; and Schwaighofer, A. 2012. Budget optimization for sponsored search: Censored learning in MDPs. *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence* 54–63.

Antos, A., and Kontoyiannis, I. 2001. Convergence properties of functional estimates for discrete distributions. *Random Structures and Algorithms* 19(3-44):163–193.

Audibert, J.-Y.; Bubeck, S.; and Munos, R. 2010a. Best arm identification in multi-armed bandits. In *COLT*.

Audibert, J.; Bubeck, S.; and Munos, R. 2010b. Best arm identification in multi-armed bandits. *In Proceedings of the 23rd Annual Conference on Learning Theory (COLT)* 41–53.

Audibert, J.-Y.; Munos, R.; and Szepesvári, C. 2009. Exploration-exploitation trade-off using variance estimates in multi-armed bandits. *Theor. Comput. Sci.* 410(19):1876–1902.

Avner, O.; Mannor, S.; and Shamir, O. 2012. Decoupling exploration and exploitation in multi-armed bandits. *Proceedings of the 29th International Conference on Machine Learning (ICML-12)* 409–416.

Berger, A. L.; Pietra, V. J. D.; and Pietra, S. A. D. 1996. A maximum entropy approach to natural language processing. *Computational Linguistics* 22(1):39–71.

Brown, D. B. 2007. Large deviations bounds for estimating conditional value-at-risk. *Operations Research Letters* 35(6):722–730.

Bubeck, S.; Munos, R.; and Stoltz, G. 2011. Pure exploration in finitely-armed and continuously-armed bandits. *Theor. Comput. Sci.* 412:1832–1852.

Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, Learning, and Games*. Cambridge University Press.

Chancelier, J.-P.; Lara, M. D.; and de Palma, A. 2009. Risk aversion in expected intertemporal discounted utilities bandit problems. *Theory and Decision* 67(4):433–440.

David, H. A., and Nagaraja, H. N. 2003. *Order Statistics*. Wiley-Interscience, 3rd edition.

Denardo, E. V.; Park, H.; and Rothblum, U. G. 2007. Risk-sensitive and risk-neutral multiarmed bandits. *Math. Oper. Res.* 32(2):374–394.

Domingo, C.; Gavaldá, R.; and Watanabe, O. 2002. Adaptive sampling methods for scaling up knowledge discovery algorithms. *Data Mining and Knowledge Discovery* 6(2):131–152.

Even-Dar, E.; Mannor, S.; and Mansour, Y. 2002. PAC bounds for multi-armed bandit and Markov decision processes. In *COLT*.

Gabillon, V.; Ghavamzadeh, M.; Lazaric, A.; and Bubeck, S. 2012. Best arm identification: A unified approach to fixed budget and fixed confidence. In *NIPS*.

Gittins, J.; Glazebrook, K.; and Weber, R. 2011. *Multi-armed Bandit Allocation Indices*. Wiley.

Kalyana, S.; Tewari, A.; Auer, P.; and Stone, P. 2012. PAC subset selection in stochastic multi-armed bandits. In *ICML*.

Karnin, Z.; Koren, T.; and Somekh, O. 2013. Almost optimal exploration in multi-armed bandits. *Proceedings of the 30th International Conference on Machine Learning* 7:1238–1246.

Lai, T. L., and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6:4–22.

Mannor, S., and Tsitsiklis, J. N. 2011. Mean-variance optimization in markov decision processes. In *ICML*.

Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7:77–98.

Maron, O., and Moore, A. W. 1993. Hoeffding races: Accelerating model selection search for classification and function approximation. *Proceedings of the Seventh Annual Conference on Neural Information Processing Systems* 59–66.

Mnih, V.; Szepesvári, C.; and Audibert, J. 2008. Empirical bernstein stopping. *Proceedings of the 25th International Conference on Machine Learning* 672–679.

Pál, D.; Poczós, B.; and Szepesvári, C. 2010. Estimation of Rényi entropy and mutual information based on generalized nearest-neighbor graphs. *Advances in Neural Information Processing Systems 23* 1849–1857.

Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin of the AMS* 55:527–535.

Sani, A.; Lazaric, A.; and Munos, R. 2012. Risk-aversion in multi-armed bandits. In *ICML Workshop*.

Schied, A. 2006. Risk measures and robust optimization problems. *Stoch. Models* 22(4):753–831.

Sricharan, K.; Raich, R.; and Hero, A. 2011. k-nearest neighbor estimation of entropies with confidence. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, 1205–1209.

Yu, J. Y., and Nikolova, E. 2013. Sample complexity of risk-averse bandit-arm selection. In *Proceedings of IJCAI*.

Appendix: Proofs

Proof of Lemma 1. Since $\widehat{G}_i$ are proper estimators, we have:

$$\begin{aligned}\mathbb{P}(\widehat{G}_i(n) \geqslant \widehat{G}_{i^*}(n)) &\leqslant \mathbb{P}\Big(\widehat{G}_i(n) > G_i + \frac{\gamma_i}{2}\Big) \\ &+ \mathbb{P}\Big(\widehat{G}_{(i^*)}(n) < G_{(i^*)} - \frac{\gamma_i}{2}\Big) \\ &\leqslant 2Q\Big(n, \frac{\gamma_j}{2}\Big)\end{aligned}$$

□

Proof of Lemma 2. Let $e_p(l)$ denote the probability that the best arm is not eliminated until episode l but will be dropped out at this episode. Consider the $|S_l| - 1$ suboptimal arms, and suppose that $i(1), i(2), \ldots, i(|S_l| - x_l)$ are the best $|S_l| - x_l$ arms among them. The best arm is eliminated if the estimate of its functional is lower than all of arms $i(1), i(2), \ldots, i(|S_l| - x_l)$. Given this, we have:

$$\begin{aligned}e_p(l) &\leqslant \sum_{j=1}^{|S_l|-x_l} P(\widehat{G}_{i^*}(n_l) \leqslant \widehat{G}_{i(j)}(n_l)) \\ &\leqslant 2(|S_l| - x_l)Q\Big(l\frac{T-H}{H}, \frac{d}{2}\Big)\end{aligned}$$

where n_l denotes the total number of pulls per arm i up to round l. Since $n_l = l\Big\lfloor\frac{T}{H}\Big\rfloor \geqslant l\Big(\frac{T}{H} - 1\Big)$, we obtain the desired inequality by replacing n_l with $l\frac{T-H}{H}$ (recall the function Q is strictly monotone decreasing in n). □

Proof of Theorem 7. Recall that $\widehat{\mu}(j)$ and $\widehat{\sigma}^2(j)$ are unbiased estimators. By Hoeffding's inequality, we have

$$\mathbb{P}(|\widehat{\mu}(j) - \mu(j)| > \gamma_j/4) \leqslant 2\exp\left(-\frac{N\gamma_j^2}{8(B-A)^2}\right).$$

Observe that $(X_k^i - \widehat{\mu}(j))^2 \in [0, (B-A)^2]$ with probability 1. By the triangle inequality, Hoeffding's inequality, we have

$$\begin{aligned}&\mathbb{P}\left(\left|\widehat{\sigma}^2(j) - \sigma^2(j)\right| > \gamma_j/(4\lambda)\right) \\ &\leqslant 2\exp\left(-\frac{N(\frac{N-1}{N}\gamma_j/\lambda)^2}{8(B-A)^4}\right).\end{aligned}$$

Observe that

$$\begin{aligned}&\{\left|-\widehat{\mu}(j) + \lambda\widehat{\sigma}^2(j) + \mu(j) - \lambda\sigma^2(j)\right| > \varepsilon\} \\ &\subseteq \{|\widehat{\mu}(j) - \mu(j)| + \left|\lambda\widehat{\sigma}^2(j) - \lambda\sigma^2(j)\right| > \varepsilon\},\end{aligned}$$

Hence, we have

$$\begin{aligned}
&\mathbb{P}\left(\left|\widehat{G}_i^{M,\lambda}(N) - G_i^{M,\lambda}\right| > \gamma_j/2\right) && (2)\\
&= \mathbb{P}(\left|-\widehat{\mu}(j) + \lambda\widehat{\sigma}^2(j) + \mu(j) - \lambda\sigma^2(j)\right| > \gamma_j/2) && (3)\\
&\leqslant \mathbb{P}(|\widehat{\mu}(j) - \mu(j)| + \left|\lambda\widehat{\sigma}^2(j) - \lambda\sigma^2(j)\right| > \gamma_j/2) && (4)\\
\text{(Union bound)}\quad &\leqslant \mathbb{P}(|\widehat{\mu}(j) - \mu(j)| > \gamma_j/4) && (5)\\
&+ \mathbb{P}(\lambda\left|\widehat{\sigma}^2(j) - \sigma^2(j)\right| > \gamma_j/4), && (6)
\end{aligned}$$

where the last inequality follows from a union bound.

Finally, we have

$$\begin{aligned}
&\mathbb{P}(\widehat{G}_i^M(N) - \widehat{G}_{i^*}^M(N) \geqslant 0)\\
&= \mathbb{P}\Big(\widehat{G}_i^M(N) - G_i^M + G_i^M -\\
&G_{i^*}^M + G_{i^*}^M - \widehat{G}_{i^*}^M(N) \geqslant 0\Big)\\
&= \mathbb{P}\Big(\widehat{G}_i^M(N) - G_i^M + G_{i^*}^M - \widehat{G}_{i^*}^M(N) \geqslant \gamma_j\Big)\\
\text{(Union bound)}\quad &\leqslant \mathbb{P}(\widehat{G}_i^M(N) - G_i^M \geqslant \gamma_j/2)\\
&+ \mathbb{P}(G_{i^*}^M - \widehat{G}_{i^*}^M(N) \geqslant \gamma_j/2)\\
&\leqslant \mathbb{P}(\left|\widehat{G}_i^M(N) - G_i^M\right| \geqslant \gamma_j/2)\\
&+ \mathbb{P}(\left|G_{i^*}^M - \widehat{G}_{i^*}^M(N)\right| \geqslant \gamma_j/2)\\
\text{(by (6))}\quad &\leqslant 2\exp\left(-\frac{N\gamma_j^2}{8(B-A)^2}\right)\\
&+ 2\exp\left(-\frac{N(\frac{N-1}{N}\gamma_j/\lambda)^2}{8(B-A)^4}\right).
\end{aligned}$$

□

Proof of Theorem 10. By Theorem 9, we have

$$\left|\mathbb{E}X^i_{(\lceil\lambda N\rceil)} - q_i(\lambda)\right| = \left|V_N^i\right|. \tag{7}$$

By the Triangle Inequality, Equation (7), and Chebyshev's Inequality, we have

$$\begin{aligned}
&\mathbb{P}\left(\left|\widehat{G}_i^{V,\lambda}(N) - G_i^{V,\lambda}\right| \geqslant \varepsilon\right) \\
&= \mathbb{P}\left(\left|X^i_{(\lceil \lambda N \rceil)} - q_i(\lambda)\right| \geqslant \varepsilon\right) \\
&\leqslant \mathbb{P}\left(\left|X^i_{(\lceil \lambda N \rceil)} - \mathbb{E}\widehat{G}_i^{V,\lambda}(N)\right| + \left|\mathbb{E}X^i_{(\lceil \lambda N \rceil)} - q_i(\lambda)\right| \geqslant \varepsilon\right) \\
&\leqslant \mathbb{P}\left(\left|X^i_{(\lceil \lambda N \rceil)} - \mathbb{E}X^i_{(\lceil \lambda N \rceil)}\right| \geqslant \varepsilon - \left|V_N^i\right|\right) \\
&\leqslant \frac{\mathbb{V}\, X^i_{(\lceil \lambda N \rceil)}}{(\varepsilon - \left|V_N^i\right|)^2} \leqslant \frac{W_N^i}{(\varepsilon - \left|V_N^i\right|)^2}.
\end{aligned}$$

□

Proof of Lemma 11. We have:

$$\begin{aligned}
\mathbb{P}(\widehat{G}_i^{V,\lambda}(N) \geqslant \widehat{G}_{i^*}^{V,\lambda}(N)) \quad &\leqslant \quad \mathbb{P}(\widehat{G}_i^{V,\lambda}(N) \geqslant G_i^{V,\lambda} + \gamma_j/2) \\
&+ \quad \mathbb{P}(\widehat{G}_{i^*}^{V,\lambda}(N) \leqslant G_{i^*}^{V,\lambda} - \gamma_j/2) \\
&\leqslant \quad \mathbb{P}(\widehat{G}_i^{V,\lambda}(N) \geqslant G_i^{V,\lambda} + d/2) \\
&+ \quad \mathbb{P}(\widehat{G}_{i^*}^{V,\lambda}(N) \leqslant G_{i^*}^{V,\lambda} - d/2) \\
&\leqslant \quad \mathbb{P}(|\widehat{G}_i^{V,\lambda}(N) - G_i^{V,\lambda}| \geqslant d/2) \\
&+ \quad \mathbb{P}(|\widehat{G}_{i^*}^{V,\lambda}(N) \leqslant G_{i^*}^{V,\lambda}| \geqslant d/2) \\
&\leqslant \quad 2\frac{W_N^i}{(\frac{d}{2} - \left|V_N^i\right|)^2}
\end{aligned}$$

□

Proof of Theorem 13. By the definitions of ρ_λ^A and $\widehat{Y}_\lambda^i$, and by the Triangle

Inequality, we have

$$\begin{aligned}
&\lambda \left|\rho_\lambda^A(X^i) - (-\widehat{Y}_\lambda^i)\right| = \\
&\left|\int_0^\lambda q_i(\xi)\, d\xi - \left(\sum_{j=0}^{\lfloor \lambda N\rfloor -1} \frac{X^i_{(j+1)}}{N} + \left(\lambda - \frac{\lfloor \lambda N\rfloor}{N}\right) X^i_{(\lceil \lambda N\rceil)}\right)\right| \\
&\leqslant \left|\int_0^{\lambda'} q_i(\xi)\, d\xi - \left(\sum_{j=0}^{\lambda' N-1} \frac{1}{N} X^i_{(j+1)}\right)\right| \\
&\leqslant \underbrace{\left|\int_0^{\lambda'} q_i(\xi)\, d\xi - \left(\sum_{j=0}^{\lambda' N-1} \frac{1}{N} \mathbb{E}X^i_{(j+1)}\right)\right|}_{R} \\
&+ \underbrace{\left|\left(\sum_{j=0}^{\lambda' N-1} \frac{1}{N} \mathbb{E}X^i_{(j+1)}\right) - \left(\sum_{j=0}^{\lambda' N-1} \frac{1}{N} X^i_{(j+1)}\right)\right|}_{S}.
\end{aligned}$$

Observe that by Theorem 9, we have

$$\begin{aligned}
R &\leqslant \left|\int_0^{\lambda'} V_N^i(\lambda) \mathrm{d}\lambda\right| \\
&\leqslant \int_0^{\lambda'} \left(\frac{\lambda(1-\lambda)D'}{2D^3(N+2)} + \frac{C_1}{N^2}\right) \mathrm{d}\lambda \\
&= \frac{(\lambda'^2/2 - \lambda'^3/3)D'}{2D^3(N+2)} + \frac{C_1\lambda'}{N^2} \triangleq Q.
\end{aligned}$$

By assumption on N, we can verify by simple algebra that $Q < \lambda\varepsilon/2$.

Observe that

$$\begin{aligned}
&\mathbb{P}\left(\left|\rho_\lambda^A(X^i) - (-\widehat{Y}_\lambda^i)\right| > \varepsilon\right) \\
&\leqslant \mathbb{P}(R + S > \lambda\varepsilon) \\
&= \mathbb{P}(S > \lambda\varepsilon - R) \\
&\leqslant \mathbb{P}(S > \lambda\varepsilon - Q) \\
&= \mathbb{P}\left(\left|\sum_{j=0}^{\lambda' N-1} \frac{1}{N}\mathbb{E}X^i_{(j+1)} - \frac{1}{N}X^i_{(j+1)}\right| > \lambda\varepsilon - Q\right) \\
&\leqslant 2\exp\left(-\frac{(\lambda\varepsilon - Q)^2 N^2}{2\lambda' N(2M)^2}\right) \leqslant 2\exp\left(-\frac{(\lambda\varepsilon/2)^2 N}{8\lambda' M^2}\right) \leqslant \delta.
\end{aligned}$$

where the last two inequalities follow by Azuma's Inequality for bounded-difference martingale sequences and the assumption on N. □